arXiv:1402.3305v1 [cs.DC] 11 Feb 2014

Real-Time Notification for Resource Synchronization

Martin Klein[1], Robert Sanderson[1], Herbert Van de Sompel[1], and Michael L. Nelson[2]

[1] Los Alamos National Laboratory, Los Alamos USA
{mklein, rsanderson, herbertv}@lanl.gov
[2] Old Dominion University, Norfolk USA
mln@cs.odu.edu

Abstract. Web applications frequently leverage resources made available by remote web servers. As resources are created, updated, deleted, or moved, these applications face challenges to remain in lockstep with the server's change dynamics. Several approaches exist to help meet this challenge for use cases where "good enough" synchronization is acceptable. But when strict resource coverage or low synchronization latency is required, commonly accepted Web-based solutions remain elusive. This paper details characteristics of an approach that aims at decreasing synchronization latency while maintaining desired levels of accuracy. The approach builds on pushing change notifications and pulling changed resources and it is explored with an experiment based on a DBpedia Live instance.

1 Introduction

The Web is highly dynamic [4,6,11], with resources continuously being created, updated, deleted, and moved. Digital libraries that leverage third party resources face the challenge of keeping in step with this rate of change. There are significant use cases in digital libraries that require real-time and accurate synchronization, and in many cases this need is addressed through ad-hoc technical approaches implemented by a small group of collaborating systems. Proposals for more generic, Web-scale, synchronization approaches have been suggested but have not been widely adopted.

ResourceSync, a project launched by the National Information Standardization Organization[3] (NISO) and the Open Archives Initiative[4] (OAI) aims to design an approach for resource synchronization that is aligned with the Web Architecture [7] and that has a fair chance of adoption by different communities. The ResourceSync effort recognizes the challenge of devising an approach that can be applied across a variety of use cases that entail different types of resources, different types of changes, and differing requirements regarding the

[3] http://www.niso.org
[4] http://www.openarchives.org

coverage and speed of synchronization. However, the current specification [9] is focused on a pull-based synchronization approach and does not consider push-based implementations.

In this paper we present a design for a push-based approach as an extension of the Resourcesync framework and investigate its performance for a synchronization use case that entails resources that change at a very high frequency. We synchronize two systems with an instance of DBpedia Live [1] via our prototype push-based architecture and measure synchronization latency and accuracy. This use case is representative for modern digital libraries that rely on Linked Data to integrate constantly changing resources from various datasets. For example, the BBC Linked Data applications integrate data from Last.FM, DBpedia, MusicBrainz, GeoNames, and others.

2 Related Work

Resource synchronization is not a new problem. It has been acknowledged as an issue, among others, by Tim Berners-Lee [2], Umbrich et al. [19], and the related W3C Dataset Dynamics activity[5]. The herein presented work is embedded in the ResourceSync project, which has published a specification for Web resource synchronization [9] (a beta draft at the time of writing). The theoretical foundations of the ResourceSync framework were detailed in [20] and its technical aspects were described in [8]. The below overview of work in the problem domain distinguishes between approaches that rely on optimizing pull interactions with resources, those that are based on communicating change notifications, and others that focus on the transfer of changed content only.

2.1 Push/Pull Approaches

Synchronization approaches are based on a pull method, a push method, or a hybrid of both. The issue of which approach to use under which circumstances has been the subject of several research endeavors. For example, Bhide et al. [3] theoretically compare push and pull approaches for disseminating dynamic web data with an emphasis on a client's temporal coherency requirement. They introduce *Push and Pull (PaP)*, a method with a scheduled pull being the default mechanism but where the server also has the capability to push changes if it foresees that the polling client would miss the changes otherwise. Another concept they discuss is *Push or Pull (PuP)* where the push method is the default and the server can dynamically allocate push or a pull channels to clients, depending on available resources. Silberstein et al. [17] explored the boundary conditions where it was optimal to push or pull feeds, depending on frequency of reads and writes. Their work suggests that in order to setup a resource-aware synchronization implementation local decisions on a creator/consumer basis are necessary. In this sense, they show that the push method is preferable if the

[5] `http://www.w3.org/wiki/DatasetDynamics`

user's consumption frequency is greater than the event creation frequency and pull otherwise.

2.2 Change Notifications

The Event Notification Protocol [13] specifies requirements for change notifications in relation to WebDAV implementations. *DSNotify* introduced by Popitsch and Haslhofer [12] is a change detection and notification framework for Linked Datasets. Volz et al. [21] introduced *Silk*, a linking framework based on the Web of Data Link Maintenance Protocol[6]. Ping the Semantic Web[7] offers change notification as a web service. All three systems are geared towards Linked Datasets and are not designed as generic resource synchronization frameworks. They further rely on aggregated baseline data to be available in a central service, raising scalability concerns in light of an ever expanding Linked Data cloud. *Pingback* and its extension *Semantic Pingback* [18] both provide a lightweight notification approach. However, the approaches are based on subscriptions to individual resources, making them problematic for large resource collections. The HTTP-based Simple Update Protocol (SUP)[8] and the UDP-based Simple Lightweight Announcement Protocol (SLAP)[9] provide conceptual mechanisms to notify about change events. However, they both suffer from the lack of acceptance and reference implementations.

3 Architectural Paradigms

At the core of the resource synchronization problem is the need for one or more Destination servers to remain synchronized with (some of the) resources made available by a Source server. For typical synchronization scenarios we can distinguish between a change notification (CN) requirement, which allows a Destination to understand that a resource has changed at the Source and what the type of change is (e.g., update, delete) and a content transfer (CT) requirement, which allows a Destination to update its holdings to reflect the change the resource underwent at the Source.

Pull-based approaches are inarguably the simplest alternative for resource synchronization as Destinations recurrently poll the Source for changes and retrieve resources that have to be synchronized. However, these approaches are not subject of this paper. They have been explored in our previous work [20] and implemented in the ResourceSync framework [9]. Therefore, in this section we consider two push-based synchronization architectures that adhere to the separation between CN and CT.

[6] `http://www4.wiwiss.fu-berlin.de/bizer/silk/wodlmp/`
[7] `http://pingthesemanticweb.com/`
[8] `http://code.google.com/p/simpleupdateprotocol/`
[9] `http://joecascio.net/joecblog/2009/05/18/announcing-slap/`

3.1 CN Simulated Push & CT Pull

Introducing feed-based technology (e.g., Atom/RSS, OAI-PMH repository) at the Source was a first step to improve pull-based approaches. This technology introduces a new component into the synchronization architecture other than Source and Destination: a Service. More recent approaches, such as PubSubHubbub (PuSH), have decoupled such a Service from the Source, reducing the memory burden the latter carries but adding interactions. Such an architecture is displayed in 1 of Figure 1.

Typical PuSH use cases are concerned with the transfer of Atom records, commonly conveying news items, from publishers to subscribers. But PuSH could also be applied to transfer CNs in which case a feed is maintained by the Source. When the Source's feed content is updated (i.e., one or more resources changed), the Source sends a content-less ping to a Hub Service (CN interaction 1.1). The Hub Service then retrieves the Atom feed from a previously registered location (CN interaction 1.2) and consecutively uses AtomPub to push Atom entries to Destinations (CN interaction 1.3). In order to make this possible, Destinations have previously subscribed to the Hub Service and have provided a callback URI. Upon receipt, the Destination pulls changed resources, if required (CT interaction 1.4). Since interactions 1.1 and 1.2 are a ping and a pull (but not a push), we refer to this approach as simulated push.

From the perspective of the Source, this architecture is convenient as it moves the burden of maintaining a longer-term change notification memory to the Hub Service. The Hub maintains a feed per Source and, per feed, keeps track of each Destination's callback URI. The initial ping is easy to implement for the Source, and reduces unnecessary interactions. However, without the introduction of special-purpose optimizations the Hub will repeatedly retrieve the same CN entries. The requirement in PuSH for the Destination to have a web presence for callback may not be appropriate or even possible in certain use cases, for example when the Destination resides on a mobile device.

A scenario we call selective synchronization, refers to a Destination's need to only remain in sync with a subset of resources made available by a Source, or to only be aware of certain types of its changes. The notion of a channel as a conduit for notifications about a subset of a Source's changes follows quite naturally. An example could be a Source defining channels to reflect DBpedia categories. However, for selective synchronization, simulated push is problematic. The Source would have expose multiple feeds and the Hub Service must maintain each of them. Also, it is likely that Destinations have to maintain one callback per channel they are interested in. If the Source has many such channels, and if resources are represented in multiple channels, the number of interactions involved (CN interactions 1.1 and 1.2) increases significantly.

3.2 CN Push & CT Pull

A real, rather than a simulated, push approach for CNs (2 in Figure 1) is the focus of this paper. It consists of the Source pushing CNs via the intermediation

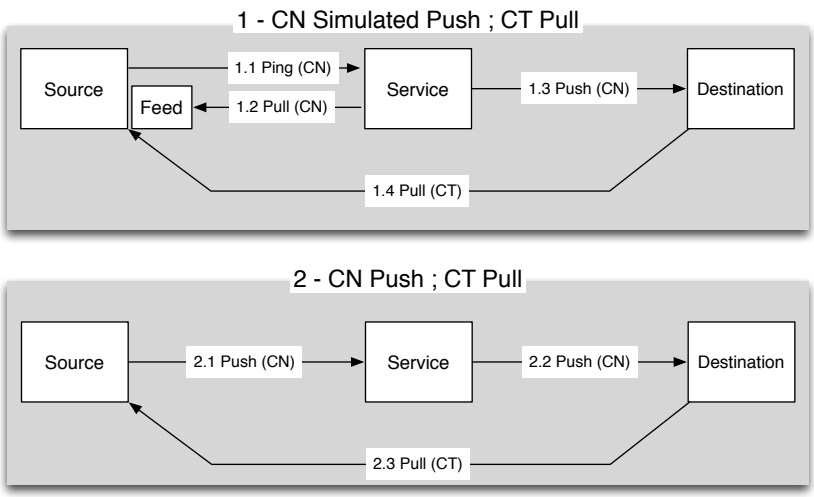

Fig. 1. Architectures Paradigms

of a Service to Destinations (CN interactions 2.1 and 2.2). As in the previous approaches, upon receipt of a notification, the Destination may decide to obtain the changed resource (CT interaction 2.3).

Interestingly, the memory requirements for the Source are zero: it can fire off CNs and immediately forget about them. The number of transactions is reduced as the initial ping of the previous architecture is no longer required, but the real optimization comes in terms of the Service not having to pull the Source's feed(s) for every change. The payload for the CNs could be designed to meet synchronization requirements, rather than conveying them in a less than suitable container like Atom entries.

An example technology in this space is XMPP [15,16], used for instant messaging between humans but also for machine to machine communication [14]. This technology requires the Destination to listen when notifications are transmitted; this is not dissimilar to being online for the callback method of PuSH. XMPP has its own protocol control structure in XML and is approximately the same order of magnitude in bytes as the Atom structure, however redundant elements are not passed backwards and forwards reducing required bandwidth.

The XMPP PubSub extension [10] provides supports for channels, with a Service rather than both Service and Source in charge of maintaining them. The Source publishes the messages to appropriate channels, which it can create as desired. The Service passes these messages on to Destinations that subscribe to the channel. XMPP also has an extension that allows the nesting of channels so that a subscriber to a channel receives messages from all of its child channels. For a large and dynamic set of resources, this capability significantly reduces unnecessary transactions.

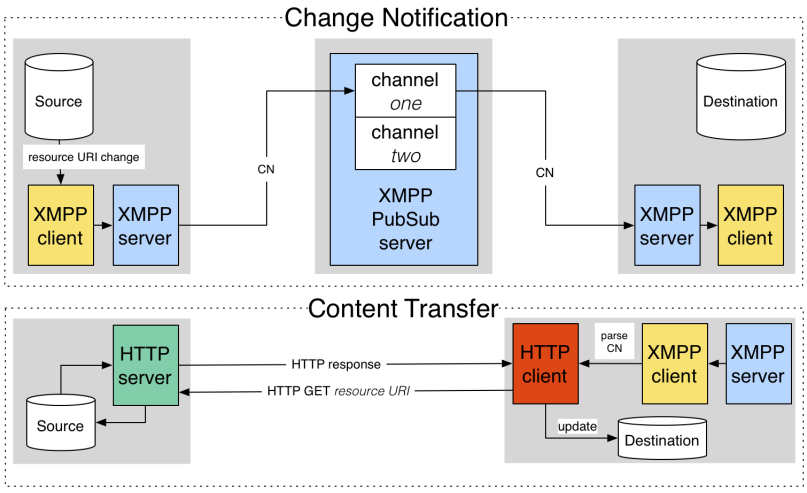

Fig. 2. An architecture based on XMPP PubSub

4 Experiment

In order to explore the problem domain, an example resource synchronization approach was formulated and tested. The approach consists of pushing CNs from a Source to Destinations. The format and content of the CNs are not subject of this experiment. For a detailed discussion, please refer to our previous work [8]. CT consists of a Destination pulling the entire resource about which a Destination received a change notification. The approach was tested to synchronize two Destinations with a Source that is a DBpedia Live instance operated at the Los Alamos National Laboratory (LANL).

There are several motivations for these choices. First, little is known regarding reliability and scalability of change notification approaches, as recognized in [19]. Also, the considerable update frequency of DBPedia Live topic-URI graphs, which was observed to average around two changed graphs per second during a two month period, provides a credible synchronization challenge. Finally, intuition suggests that such high change frequency may cause problems, especially for an approach where entire resources rather than only changes are communicated.

4.1 Synchronization Approach

The tested approach (Figure 2) handles CNs by means of a push paradigm. XMPP PubSub is used as the protocol to push notifications from Source to Destinations. Consequently, Source and Destination each have a Jabber ID, operate an XMPP client to send/receive notifications, and have an XMPP home server. XMPP PubSub nodes serve as channels for selective synchronization. The Source creates the channels and publishes notifications to them. Destinations subscribe

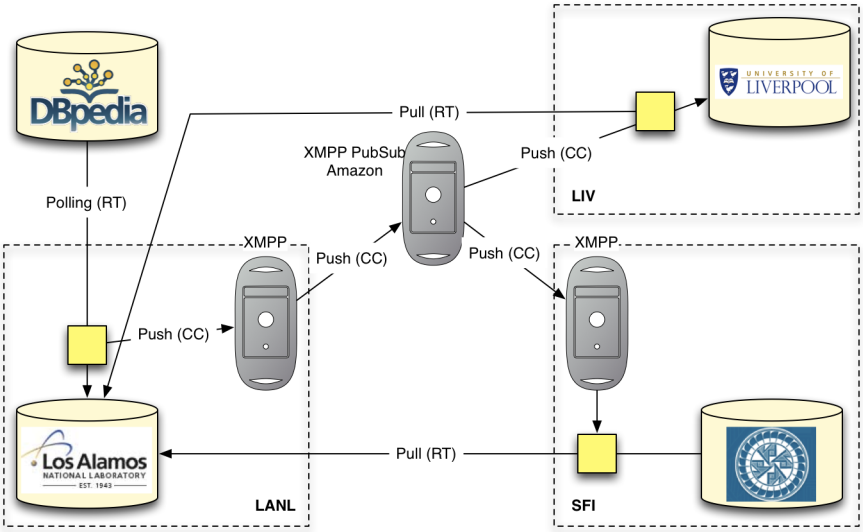

Fig. 3. Systems involved in the experiment

to channels and as a result receive notifications. The XMPP PubSub server manages channels and subscriptions, and relays CNs from Source to Destinations. CT is handled by means of a pull of the entire changed resource. When a Destination receives a CN, it determines whether it should obtain the corresponding updated version of the resource. If so, it dereferences that resource's URI. Note that, when a resource changes very rapidly or when the Destination processes CNs with a delay, the possibility arises that the resource version pulled by the Destination is more recent than the event time reported in the CN.

4.2 Configuration

To create a baseline for synchronization, the LANL DBpedia Live instance (bottom left of Figure 3) uploaded a dump from DBpedia Live (top left of Figure 3). Since then, it has kept in step with changes using the latter's content-type-specific changeset mechanism[10] that consists of providing an updated and deleted changeset file per update cycle, in which each line is a triple that has a DBpedia Live topic-URI as subject. The LANL instance continuously monitors the availability of new changesets and, when available, processes them. The heuristics involved in doing so are beyond the scope of this paper. It suffices to say that at the end of each processing cycle, the LANL instance is updated through the addition of new versions of previously existing topic-URI graphs, the deletion of topic-URI graphs, or the creation of new ones. Each of these events is communicated on the `dbpedia_all` channel that covers all changes to all DBpedia

[10] See `http://live.dbpedia.org/liveupdates/`

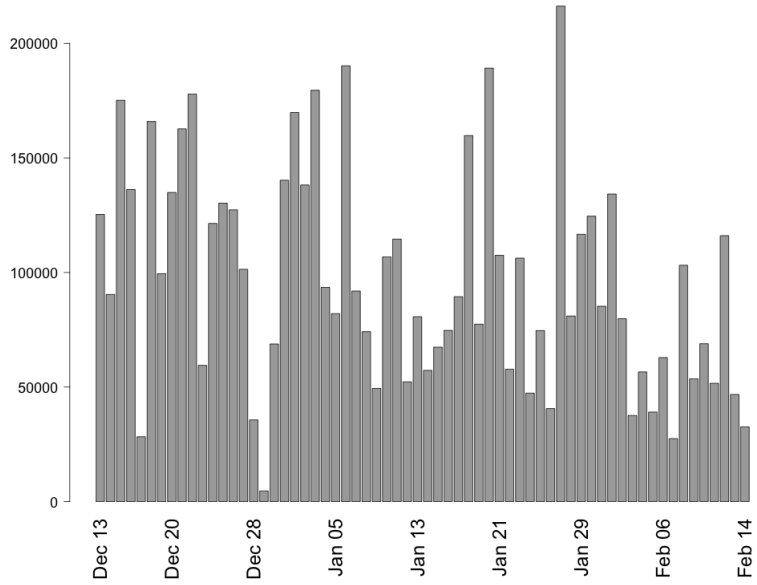

Fig. 4. CNs on dbpedia_all per day

topic-URIs. Figure 4 shows the resulting daily traffic. Overall, 99% of these CNs are about updates, 0.6% about deletions, and 0.03% about creations. In addition, depending on the DBpedia category of a topic-URI, CNs are also sent to channels such as `dbpedia_music`, `dbpedia_business`, etc. The LANL instance holds on to all versions of topic-URI graphs and stores its serializations.

An XMPP PubSub server operated in Amazon's cloud (center of Figure 3) provides the infrastructure for the channels. The LANL instance creates the channels to which Destinations can subscribe. In the experiment, two remote servers subscribe. A server at the University of Liverpool (LIV, top right of Figure 3) subscribes to dbpedia_all. Upon receipt of the CNs, it updates its local DBpedia collection by dereferencing or deleting the corresponding topic-URIs. It stores serialized graphs in its file system in the same way the LANL instance does. A server at the Santa Fe Institute (SFI, bottom right of Figure 3) remains synchronized in the same way, but its topic-URI collection is maintained in a Virtuoso triple store. Neither of these Destinations performed a baseline synchronization and they do not hold on to old versions. Each runs an XMPP client that connects to a local XMPP server.

Figure 3 depicts the information flow in the experimental setup: the LANL instance is polling DBpedia Live for CT; the LANL instance sends CNs about changes it is undergoing and these CNs are relayed by the XMPP PubSub server in the cloud to the Liverpool and Santa Fe servers. These servers then use HTTP GET to obtain new or updated resources (CT).

Run	Total CNs	Diff LANL-LIV	MaxQ LIV	MaxQ SFI
1	13,819	4 (0.03%)	120	57
2	32,453	2 (0.006%)	389	65
3	6,910	1 (0.01%)	93	44
4	24,400	5 (0.02%)	175	77
5	11,850	0 (0.0%)	84	57
6	14,937	4 (0.03%)	98	61
total	104,369	16 (0.015%)		

Table 1. Experimental Results

4.3 Results

A first experiment investigated whether the Destinations at LIV and SFI were able to keep up with the pace of CNs sent out by LANL. As CNs arrived at the Destinations, they were pushed into a first-in-first-out queue awaiting processing. This entailed obtaining the changed resource from LANL in case of create/update events and updating the local versions for all events. The size of this queue as compared to the amount of received CNs serves as an indication of the capability of a Destination to remain in lockstep. These parameters were observed during six eight-hour runs conducted at different times of the day, in the course of a single week. During these sessions, the output rate of changesets from DBpedia Live was very inconsistent. As a result, LANL sometimes broadcast thousands of CNs in a single five minute interval, and at other times not a single one for several hours. Even in less extreme situations, CNs were sent out in a bursty manner as a result of LANL polling DBpedia Live every 30 seconds for changesets but pushing out all resulting notifications once those were processed. Due to this bursty sending, the queues at the Destinations could grow rapidly. Nevertheless, both Destinations were able to keep up with the barrage of CNs. The second column in Table 1 shows the total number of CNs processed per run, whereas the "MaxQ" columns show the maximum queue sizes for each run for LIV and SFI, respectively. As can be observed, LIV had significantly larger queues than SFI; this is due to additional time required to dereference resources from a location on another continent. Figure 5 depicts the number of CNs over the 8 hour session of the fourth run measured at five minute intervals, as well as the maximum queue size during each interval. It needs to be noted that for all runs, both LIV and SFI ended up with an empty queue at the end of each of these five minute intervals.

A second experiment investigated whether resources were effectively synchronized. This experiment was conducted during the same six eight-hour runs. It involved LANL as the Source and LIV as the Destination because both use the same storage mechanism for topic-URI graphs, allowing for a byte-level comparison of the respective collections. Lack of control over the incoming data from DBpedia Live caused a problem with dereferencing and hence synchronizing some of the communicated URIs. For example, some DBpedia Live URIs con-

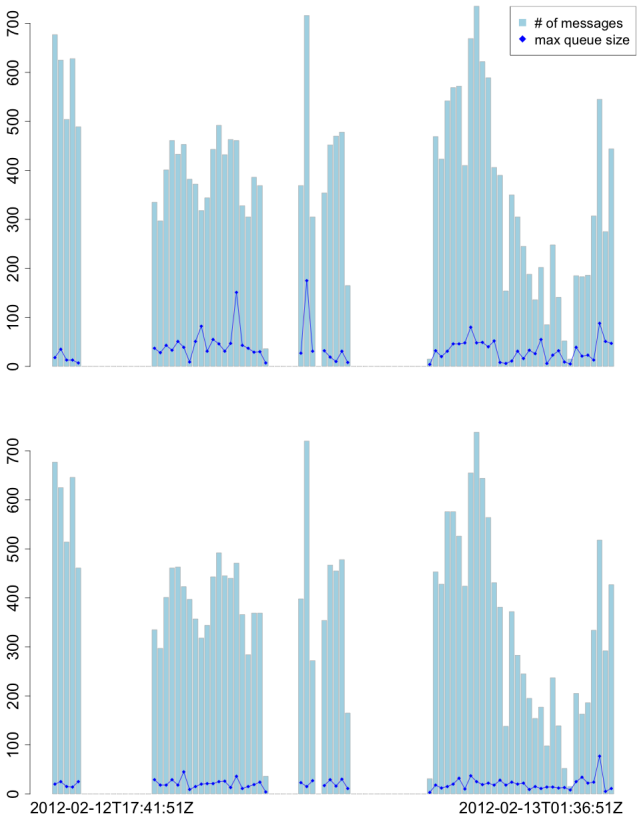

Fig. 5. Run 4, LIV (top), SFI (bottom) Destinations

tained the same character both percent-encoded and unencoded, and the LANL storage layer was unable to resolve such identifiers successfully. Occurrences of this problem are excluded from the results, but amounted to significantly less than 0.1% of all CNs. Note that this problem results from LANL having to mint and communicate URIs that contain parts of the DBpedia Live in their path, and would not occur in a real-life situation. A recursive `diff` of the Source and Destination collections was used at the end of each run to detect resources that were not synchronized correctly. The "Diff LANL-LIV" column in Table 1 shows the number and percentage of resources at LIV that were out-of-sync with LANL at the end of each run. Overall, 99.99% of the synchronization was accurate after processing more than 100,000 CNs.

To verify the intuition that the changeset approach used by DBpedia Live would naturally result in a smaller content transfer payload, a third experiment was conducted. During a 24 hour period, the sizes were monitored of respectively all changesets pulled by LANL from DBpedia Live, and all HTTP GET

transactions issued from SFI to LANL that resulted from processing CNs. The changesets amounted to 149MB compressed, which corresponds to 4.2GB uncompressed. The HTTP GET transactions totaled 1.8GB uncompressed. Assuming these transactions would have had compressed response bodies, and applying a reasonable compression coefficient of about 0.1 as shown in [5], HTTP GET transactions would have totaled around 180MB.

5 Conclusions

This paper contributes to the exploration of Web-based resource synchronization approaches. It demonstrates the real possibility of using an XMPP PubSub architecture to push CNs and use a pull-based method for CT.

Clearly additional experiments are required, but results from the reported work were far beyond expectations. A more robust implementation, without the described URI encoding issues and without the need to poll DBpedia Live in order to update LANL, could likely achieve 100% synchronization without significant additional implementation effort. The extent of unnecessary messages is minimized thanks to the intermediation of Service-side channels. The synchronization latency, even under extreme load, remained very low and can further be decreased by adding additional retrieval threads. Generally, even though developers were not previously acquainted with XMPP protocols or tools, adoption was straightforward and off-the-shelf client and server tools were readily usable.

The tested approach relies on URI dereferencing for content transfer. Even though the experiment did not encounter specific problems with this regard, the following issues arose when discussing the reported work. Firstly, some resources, like those identified by hash-bang HTTP URIs[11], require client-side processing to obtain a meaningful representation. Such resources introduce additional synchronization challenges and suggest that resources to be initially considered for a synchronization framework need to be cacheable[12]. Secondly, specific representations of a resource are subject to synchronization. But the possibilities with this regard seem limited to representations that can be requested using the URI's protocol parameters. This introduces challenges related to the subjectivity of observed changes in case of personalized (e.g., geo-dependent) representations.

6 Acknowledgments

We thank Simeon Warner (Cornell University), Bernhard Haslhofer (University of Vienna) and Carl Lagoze (Michigan University) for project input and Lyudmila L. Balakireva and Harihar Shankar of the LANL Prototyping Team for implementations. Thanks to NISO's Todd Carpenter, Nettie Lagace and Peter Murray for ResourceSync support. This work is partly funded by the Sloan Foundation and the Library of Congress. Many thanks to the University of Liverpool and the Santa Fe Institute for providing computing resources.

[11] See `http://www.jenitennison.com/blog/node/154`

[12] See `http://isolani.co.uk/blog/javascript/BreakingTheWebWithHashBangs`

References

1. DBpedia Live. http://wiki.dbpedia.org/DBpediaLive.
2. T. Berners-Lee and D. Connolly. Delta: an ontology for the distribution of differences between RDF graphs, 2001. http://www.w3.org/DesignIssues/Diff.
3. M. Bhide, P. Deolasee, A. Katkar, A. Panchbudhe, K. Ramamritham, and P. Shenoy. Adaptive Push-Pull: Disseminating Dynamic Web Data. *IEEE Transactions on Computers*, 51(6):652–668, June 2002.
4. J. Cho and H. Garcia-Molina. Estimating Frequency of Change. *ACM Transactions on Internet Technology (TOIT)*, 3(3):256–290, August 2003.
5. J. D. Fernández, C. Gutierrez, and M. A. Martínez-Prieto. RDF Compression: Basic Approaches. In *Proceedings of WWW '10*, pages 1091–1092, 2010.
6. D. Fetterly, M. Manasse, M. Najork, and J. Wiener. A Large-Scale Study of the Evolution of Web Pages. *SPE*, 34(2):213–237, 2004.
7. I. Jacobs and N. Walsh. Architecture of the world wide web. http://www.w3.org/TR/webarch/.
8. M. Klein, R. Sanderson, H. Van de Sompel, S. Warner, B. Haslhofer, C. Lagoze, and M. L. Nelson. A Technical Framework for Resource Synchronization. *D-Lib Magazine*, 19(1/2), 2013.
9. M. Klein, R. Sanderson, H. Van de Sompel, S. Warner, B. Haslhofer, M. L. Nelson, and C. Lagoze. ResourceSync Framework Specification. http://www.openarchives.org/rs/resourcesync.
10. P. Millard, P. Saint-Andre, and R. Meijer. Xep-0060: Publish-subscribe. http://xmpp.org/extensions/xep-0060.html.
11. A. Ntoulas, J. Cho, and C. Olston. What's New on the Web?: The Evolution of the Web from a Search Engine Perspective. In *Proceedings of WWW '04*, pages 1–12, 2004.
12. N. Popitsch and B. Haslhofer. DSNotify: Handling Broken Links in the Web of Data. In *Proceedings of WWW '10*, pages 761–770, 2010.
13. S. Reddy and M. L. Fisher. Requirements for event notification protocol (internet draft). http://tools.ietf.org/html/draft-ietf-webdav-enpreq-01.
14. M. Rodriguez and J. Shinavier. Linked process protocol. http://xmpp.org/extensions/inbox/lop.html.
15. P. Saint-Andre. Extensible messaging and presence protocol (xmpp): Core (RFC 6120). http://tools.ietf.org/html/rfc6120.
16. P. Saint-Andre, K. Smith, and R. Tronçon. *XMPP: The Definitive Guide: Building Real-Time Applications with Jabber Technologies*. O'Reilly Media, Inc., 2009.
17. A. Silberstein, J. Terrace, B. F. Cooper, and R. Ramakrishnan. Feeding Frenzy: Selectively Materializing Users' Event Feeds. In *Proceedings of SIGMOD '10*, pages 831–842, 2010.
18. S. Tramp, P. Frischmuth, T. Ermilov, and S. Auer. Weaving a social data web with semantic pingback. In *Proceedings of EKAW'10*, pages 135–149, 2010.
19. J. Umbrich, M. Hausenblas, A. Hogan, A. Polleres, and S. Decker. Towards Dataset Dynamics: Change Frequency of Linked Open Data Sources. In *Proceedings of LDOW '10*, 2010.
20. H. Van de Sompel, R. Sanderson, M. Klein, M. L. Nelson, B. Haslhofer, S. Warner, and C. Lagoze. A Perspective on Resource Synchronization. *D-Lib Magazine*, 18(9/10), 2012.
21. J. Volz, C. Bizer, M. Gaedke, and G. Kobilarov. Discovering and Maintaining Links on the Web of Data. In *Proceedings of ISWC '09*, pages 650–665, 2009.